

**SYSTEM AND METHOD FOR COLLECTING, ASSOCIATING, NORMALIZING
AND PRESENTING PRODUCT AND VENDOR INFORMATION ON A
DISTRIBUTED NETWORK**

BACKGROUND OF THE INVENTION

5 Field Of The Invention

 The present invention relates generally to a system and method for collecting and presenting product and vendor information on a distributed network such as the Internet.

10 Background And Related Art

 It is known to sell products on a distributed network such as the Internet. Online sales or e-commerce is a rapidly growing segment of the economy. Systems for selling products on a distributed network are sometimes referred to as electronic merchandising systems or virtual storefronts. It is further known to aggregate in one user interface access to multiple online vendors to enable a user to choose among several retailers' goods. Sites containing multiple vendors are sometimes referred to as electronic or virtual malls, or shopping agents or "bots." An electronic vendor or electronic mall provides a display that generally includes images and descriptions of merchandise. These sites also generally provide the vendors' prices for the product. Shopping agents or "bots" aggregate product pricing information from multiple vendors on a single site.

25 In addition to serving as an avenue for commerce, a distributed network allows consumers to access considerable amounts of information about products. For example,

consumers can research products by accessing information provided by manufacturers, vendors, distributors, etc. Consumers also may research products through third-party sites, such as ConsumerReports.org®, that publish industry reviews of products.

Consumers may further communicate with each other to exchange product experiences and information. For example, consumers may interact on Usenet discussion groups to share information such as personal experiences with products. In addition, it has been proposed by the assignee of the present invention to survey consumers regarding the quality of particular products and/or services and to publish or advertise the results of the survey as numerical ratings. Recently, with the rapid technological advancement of the Internet, it has become further possible for individual consumers to provide narrative reviews of products and/or services, in addition to the standardized scaled ratings.

A consumer can also research information on vendors. For example, vendors typically provide on their websites information such as their shipping, billing and return policies. As with products, consumers also may communicate with other users to exchange experiences and information related to vendors on online discussion groups or at third-party sites that allow users to rate and review vendors. There further exist websites, such as gomez.com and bizrate.com, that allow users to rate vendors.

Although there is an abundance of vendor and product information on the Internet, this information is distributed over numerous websites. To access the information, consumers need to locate these various websites. However, consumers may have difficulties finding the various websites. For instance, searching under a product name on a

search engine may locate millions of websites, most of which provide little or no relevant information. Accordingly, there presently exists a need for a methodology to provide a single source for information on products and vendors.

5 Furthermore, even if a user locates the various websites containing the desired product and vendor information, the large amount of information provided is not organized for easy access by the user. Because there exists so much information, consumers may have difficulty sorting,
10 comparing and using it. Consequently, there further exists a need for a methodology to organize and present product and vendor information for easy access by consumers.

It is generally known to use a database to electronically organize and store information. In the most
15 general sense, a database is a collection of data. Various architectures have been devised to organize data in a computerized database. Typically, a computerized database includes data stored in mass storage devices, such as tape drives, magnetic hard disk drives and optical drives. The
20 three principal database architectures are termed hierarchical, network and relational. A hierarchical database assigns different data types to different levels of the hierarchy, with each record having one owner. In this way, links between data items on one level and data items on
25 a different level are simple and direct. However, a single data item can appear multiple times in a hierarchical database, which creates data redundancy. To eliminate data redundancy, a network database stores data in nodes having direct access to any other node in the database. In the
30 network database, each record has multiple owners, and there is no need to duplicate data since all nodes are universally accessible. Alternatively, in a relational database such as

Oracle®, Sybase®, Informix®, Microsoft SQL Server®, Access®, and others, the basic unit of data is a relation that comprises attributes and tuples. The records in a relational database have no owner.

5 In an implementation of a relational database, a relation corresponds to a table having rows, where each row corresponds to a tuple, and columns, where each column corresponds to an attribute. From a practical standpoint, rows represent records of related data and columns identify individual data elements. A table defining a retailer's product line may, for example, have product names, product numbers (e.g., Stock Keeping Units or SKUs), prices and other product features. Each row of this table holds data for a single product and each column holds a single attribute, such as a product name. The order in which the rows and columns appear in a table has no significance. In a relational database, one can add a new column to a table without having to modify older applications that access other columns in the table. Relational databases thus provide flexibility to accommodate changing needs.

15 All databases require a consistent structure, termed a schema, to organize and manage the information. In a relational database, the schema is a collection of tables. Similarly, for each table, there is generally one schema to which it belongs. Once the schema is designed, a tool, known as a database management system (DBMS), is used to build the database and to operate on data within the database. The DBMS stores, retrieves and modifies data associated with the database. Lastly, to the extent possible, the DBMS protects data from corruption and unauthorized access.

A human user controls the DBMS by providing a sequence of commands selected from a data sublanguage. The syntax of data sublanguages varies widely, but the American National Standards Institute (ANSI) and the International
5 Organization for Standardization (ISO) have adopted Structured English Query Language (SQL) as a standard data sublanguage for relational databases. SQL comprises a data definition language (DDL), a data manipulation language (DML) and a data control language (DCL). DDL allows users
10 to define a database, to modify its structure and to destroy it. DML provides the tools to enter, modify and extract data from the database. DCL provides tools to protect data from corruption and unauthorized access. Although SQL is standardized, most implementations of the ANSI standard have
15 subtle differences. Nonetheless, the standardization of SQL has greatly increased the utility of relational databases for many applications, including retail sales and merchandising operations.

Although access to relational databases is facilitated
20 by standard data sublanguages, users still must have detailed knowledge of the database's terminology to obtain needed information from a database since one can design many different schemas to represent the storage of a given collection of information. For example, in an electronic
25 merchandising system, a merchant may elect to store product information, such as a product SKU, product name, product description, price and tax code, within a relational database. Another merchant may elect to store a different product SKU, product name, description, price and tax code
30 in a table. In this situation, an SQL query designed to retrieve a product price from one merchant's database is not useful for retrieving the price for the same product in the

other merchant's database because the differences in data types require the use of different SQL queries. As a consequence, developers of retail applications accessing product information from relational databases have to adapt their SQL queries to each individual schema. This, in turn, prevents their applications from being used in environments where there are a wide variety of databases having different schemas, such as the World Wide Web.

The rapid development of the World Wide Web (Web) has facilitated the use of online merchant systems. Online merchant systems enable merchants to creatively display and describe their products to a global audience of shoppers using Web pages defined by an output language such as hypertext markup language (HTML). HTML enables merchants to lay out and display content, such as text, pictures, sound and video. Web shoppers access a merchant's page using a browser, such as Microsoft Explorer® or Netscape Navigator®, installed on a client connected to the Web through an online service provider, such as the Microsoft Network® or America OnLine®. The browser interprets the HTML to format and display the merchant's page for the shopper. The online merchant system likewise enables shoppers to browse through a merchant's store to identify products of interest, to obtain specific product information and to electronically purchase products after reviewing product information. Merchants often store product data, such as product descriptions, prices and pictures, in relational databases. Online merchant systems, therefore, have to interface with merchant databases to access and display product information. As each merchant organizes their product information differently, there is a large installed base of

databases having a wide variety of data types for product information.

This problem is even greater for websites that seek to advertise and sell products from a variety of online merchant systems. A problem with finding product information on the Internet is that the same product may have numerous names or identifiers depending on the merchant's site on which it is stored. In particular, a product may be identified by its model name, serial number, SKU assigned by the vendor, distributor part number, etc. Even these identifiers may vary greatly. For example, a product may have numerous model names because the name varies from country to country, the manufacturer may periodically change the product's name, or the manufacturer, consumers and merchants may use numerous different names to refer to the same product. Similarly, different vendors use different SKU numbers. As a result, a user may have great difficulty correlating product information about the same product from different sources.

Much information on products is available on the web. For example, it is well known for vendors to provide information, such as product price, on a website. U.S. Patent No. 5,740,425 by Povilus, for DATA STRUCTURE AND METHOD FOR PUBLISHING ELECTRONIC AND PRINTED PRODUCT CATALOGS, incorporated herein by reference, provides a data structure and method for creating a product database, which defines classes of product groupings and preferably includes a listing of SKUs that correspond to a product or a component of a product. The product database further includes product information for each associated SKU. Similarly, many manufacturers of products provide online information about their products. The manufacturers may

further provide technical support and assistance over the Internet. In addition, many Internet sites provide reviews of products. These sites may have writers that test and review the products. Alternatively, the sites may allow
5 users to place their opinions about a product for other users to view. These consumer-posted reviews provide special insights into products because they reflect actual experiences with the product.

However, because the product information from different
10 sources cannot be viewed together, the utility of this abundance of information is limited.

SUMMARY OF THE INVENTION

In response to these needs, the present invention provides a system and method for collecting and displaying
15 information about a product or other data object at a website server. The product information is gathered from many diverse sources, and is normalized and associated with preexisting stored information on the server. The information displayed at the website includes: (1) a general
20 description of the product; (2) a numerical user rating of the product; (3) one or more user reviews of the product; (4) one or more industry reviews of the product; (5) one or more comparisons between the product and other similar items; (6) a list of one or more vendors that sell the
25 product; (7) a list price of the product; (8) a price for the product from each of the vendors; (9) an indication of the availability of the product at each of the one or more vendors; (10) a profile on each of the one or more vendors, and (11) a rating and/or review for each of the one or more
30 vendors. The website may further provide access to discussions regarding the product and related products. The website may further suggest complementary products that may

be purchased along with the primary product being researched.

In a preferred embodiment, the website allows the user to select the product from a list of multiple products. In turn, the website may allow the user to select the list of products from a list of classes of products. Alternatively, the website may allow the user to select desired product features and then create a list of products that possess these features.

In another embodiment, the website may allow the user to add a review or rating of the product. The website may also optionally indicate what information other users have found to be useful.

In another embodiment, the website includes decision guides that suggest a product to the user in response to a user input.

Accordingly, the present invention provides a single website to provide and organize the product and vendor information available on a distributed network, such as the Internet.

According to a preferred embodiment, the invention provides three principal instrumentalities for collecting, normalizing, associating and presenting data to a user. In order to be able to carry out attribute- or parameter-based searches of a database for products or other data objects (for simplicity, hereinafter the term "product" shall be used generically to mean any data object searchable on a database, such as for example products, services, news items, demographic, historical, scientific or statistical information, financial instrument or securities information, real estate information, and the like), consistent terminology and ontology must exist in the database.

Additionally, in order to avoid having "orphaned" or non-related items of data present in the database, it is desirable to provide the capability of associating such items of data with other, similar products, based on shared attributes. Thirdly, it is desirable to reduce the time required to complete a parameter-based product search of a database.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention will be described in detail with reference to the following drawings in which:

FIG. 1 is a flowchart illustrating a method for normalizing and associating gathered product information into a database in accordance with an embodiment of the present invention;

FIG. 2 is a table for translating or normalizing diverse product identifiers to the same products to which they are referring;

FIG. 3 is a table associating core product identifiers with corresponding domains and attributes;

FIG. 4A is a database file format showing the arrangement of product information for retrieval;

FIG. 4B is a character string look up table associating a multiplicity of character strings with unique integers;

FIG. 5 is a schematic diagram of a system for collecting, storing, and outputting product information in accordance with an embodiment of the present invention; and

FIGS. 6, 7A-7C, 8 and 9 are examples of displays of information obtained as a result of the method of FIG. 1.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

In one aspect of the present invention, a method is provided for the collection and storage of product information in a database from which it can be quickly and efficiently searched by a user and the results displayed.
5 As illustrated in Fig. 1, the first step 1001 of the method is the collection of product information and associated vendor information from the Internet or from other sources.

The collecting of product and vendor information can be carried out in a variety of ways. Some of the information may already reside at a website server in association with other applications and functions. For example, a vendor's site will already contain data relating to the vendor and the products sold by the vendor. This data may be retrieved
10 by using known "scraper" technology and loaded into a database at step 1002. The data may be subsequently combined with additional information collected from other sources.

For example, the additional information may be collected manually by a human operator at step 1001 who examines various sources such as third-party websites, publications, brochures, manufacturer specification sheets, vendor advertisements, etc., for pertinent data. The human operator at step 1002 then loads this information into an information storage device such as a database contained on a
15 20 25 server. For example, the operator may examine and record the inventory and pricing information displayed on a vendor's website.

Alternatively, information may be collected directly
30 from a server controlling the third-party information source. For instance, a vendor may sell or provide a list of its inventory and the prices for the products in the

inventory in electronic form. The list then may be transferred directly from the third-party server to the information storage device.

As mentioned above, the information also may be
5 obtained automatically through the use of programs that search for desired information on a distributed network such as the Internet. Scraper programs automatically examine third-party websites and create an output forwarding desired contents of the website to the information storage device.
10 For example, a scraper program can be designed to search the website of a vendor for the prices of products sold by the vendor. The scraper may run either in real time, upon a request by the user, or in batch mode so that the vendor's prices are periodically examined and stored, such as on a
15 weekly basis. Generally, there is a different scraper program for each type of information from each information source. In this way, a scraper can be designed specifically to locate desired information on the third-party website and to interpret the format of this information.

20 The scrapers preferably create an output using Extensible Markup Language ("XML") to return information from the third-party site in a usable format. XML is a web language similar to the standard hypertext markup language ("HTML"), but the XML rules are more complex to allow more
25 varied uses. In particular, XML is more interactive and better suited for electronic commerce because the coding contains markers the simplify the standardization of information over the Internet. This feature allows the use of intelligent agents that seek out consistent information
30 and then act on what they find. Furthermore, the parsers in XML can be small and fast and can read complex hierarchical structures.

The information may be gathered through a combination of all the above methods in order to gather the information in the most efficient manner.

As the information is gathered, it is deposited into a storage device such as a database on a server for storage and easy future access. It is well known to use databases to store and organize data. For instance, the following example shows a database containing information on two vendors that sell the same product.

Example 1

Vendor's Name	Price for product	Availability	Vendor Rating	Profile
A	\$1	Yes	4.5	A.doc
B	\$2	No	4.3	B.doc

In this example, the same product is sold at Vendor A and Vendor B. Vendor A charges \$1 for the product, has the product in stock, and has a vendor rating of 4.5. The database further indicates that a profile for Vendor A is stored in the file, A.doc. Similarly, Vendor B sells the product for \$2, does not have the product in stock, has a vendor rating of 4.3, and has a profile stored in the file B.doc.

The information collected will typically contain one or more product identifiers, such as a UPC, a manufacturer model number, a distributor part number, a vendor-specific SKU, etc. The information will further include data such as the product name, type of product (domain), and various attributes of the product with specific values for each listed attribute.

In order to have the ability to perform a parameterized or even accurate search on such information, it is necessary

to have consistent and normalized data in the database. For example, a search for "XGA" will not retrieve as a "hit" data for a laptop computer in which screen size is specified as "1024x768," even though these two terms refer to the identical type of display. Accordingly, the present invention provides a normalization engine that translates or normalizes a list of attributes and values describing an object (product) into a list containing a canonical representation for each attribute and value, in addition to a canonical domain describing the product in general (such as "notebook" to describe a portable computer, which also may be identified as a "laptop" computer). For example, the domain "laptop" would be normalized to refer to the domain "notebook," where "notebook" would be selected by the data entry operator as the canonical representation. Similarly, attribute/value pairs, such as "screen_size = xga" would be normalized to "display_res = 1024x768."

This is carried out by maintaining a list of aliases or translations for canonical domains, attributes and values in the database. Each known alias for a canonical domain term, attribute term, and value term is listed in the alias list in the database with a corresponding entry identifying the canonical representation into which the alias will be translated as the object or product information is being loaded into the database. An operator may add entries by detecting new synonyms for a canonical term in an object file and indicating the canonical term for the detected synonym. All existing occurrences of the synonym term in the database are then translated into the indicated canonical term, and the synonym is then added to the alias list, such that subsequent data entries containing that

synonym will thereafter automatically be translated into the canonical representation for entry into the database.

Before the loaded information at step 1002 can be assimilated into the database, it is first determined at step 1003 whether the information pertains to an existing product already stored in the database. If so, the new information is merged into the listings for the existing product. In case of a conflict with pre-existing information for the product, a choice may be made as to which information should take precedence. If the new information can be confirmed as corresponding to updated information with respect to the stored information, then the new information may be written in place of the pre-existing information in the database. Otherwise, the pre-existing information can be selected to take precedence over the newly loaded information. Fig. 2 shows a product map or table 2000 containing a list of known product identifiers 2001, and their corresponding core product identifier 2002. The core product identifier can be an arbitrary integer selected by the operator to identify a particular product, which may be known by various identifiers, as mentioned above. In the example, both product id #2 and product id #N refer to the same core product, as indicated by the same core product identifier, 790, contained in the map.

At step 1004 it is determined whether or not the product identifier contained in the new information is found in the product map 2000. If not, at step 1005 a new product listing is created in the database with the associated attribute/value pairs for the product. When a new domain, attribute or value is added to the database it is marked as "new." New data items will not be displayed as part of a search result until an editor or operator has reviewed them

to determine their appropriate display representation, sorting order, and whether or not they can be identified as aliases for pre-existing information in the database.

If the identifier is found, at step 1006 normalization of the domains, attributes and values is initiated. It is noted that translations are performed in a product-specific manner; thus, the attribute alias list for the attribute "display_res" for a laptop does not apply to a PDA device or a desktop PC. Similarly, the value alias list for the value "1024x768" for a laptop would be specific to the attribute "display_res" within the laptop domain and would not apply to a value for an attribute. Thus, at step 1007 the domain name of the object is compared against a domain alias list, and translated into its canonical representation as indicated in the alias list. Once the canonical domain name is obtained, each of the attributes is compared with the alias list of attributes associated with the canonical domain name map at step 1008, and each value of the attribute/value pair is then compared with the canonical attribute map at step 1009. At step 1010 it is determined whether additional attribute/value pairs exist in the new information that need to be normalized. If so, the process returns to step 1008. If not, the process ends at step 1011. Alternatively, all of the attributes can be translated together at step 1008, and then all of the values associated with each attribute can be translated together at step 1009.

According to the invention, all information in the entire database can be updated to normalize data already in the database in real time as the aliases are added to the database, by maintaining the translation rules together with the data set in the database. Additionally, the

normalization process enables all attribute information to be normalized to a common unit base (e.g., normalizing all units of length into millimeters, etc.).

An example of such a domain map 3000 is shown in Fig.

5 3. Each core product identifier 3001 has a canonical domain 3002, which in turn is associated with a number of canonical attributes 3003, 3004, 3005. For each of the attributes an alias list is maintained containing all known aliases for the canonical attribute. The same applies to values for
10 each attribute. The values are sorted in numerical order where possible; for values which are not simple numbers, the sorting order can be defined by the operator on a per attribute basis. By identifying the same attribute values as pointing to the same product, it is possible to effect
15 product and domain merges in the database automatically by defining a threshold overlap level by which attributes for separate product records in the database are the same. Once the two (or more) separately stored product records have been identified as pertaining to the same product, the
20 records can be merged into a single record in the database containing all of the product attributes in one location.

The domain editor is a Java application user interface used to manipulate data in the database, such as setting the display characteristics for the domain and attribute
25 strings, allowing the operator to translate and normalize attribute and value information, editing of data values, merging attributes, and merging domains. By setting a threshold level of overlap, the normalization engine can automatically suggest to a user possible domain merges or
30 product merges.

Further, if the product information contains multiple identifiers, each of the identifiers can be compared with

the stored product identifiers, and any new identifiers may be added to the map as being associated with or mapped to the canonical representation found for at least one of the identifiers. This can be done since it is known that all the identifiers pertain to the same product, as they were bundled together in the information collected. In this way, the database can be made to "learn" new product aliases as more and more information is loaded into it, thereby associating more and more of the information stored in the database as information is added.

An association engine makes it possible to associate previously orphaned pieces of data with product records, as more aliases are added and associations made in the database.

As illustrated in FIG. 10, the present invention provides a name database 10 containing data locations 1 for storing multiple different identifiers for each of a number of products. The name database 10 may be an array with columns 20 that represent product attributes, and rows 30 that represent the different identifiers for each attribute. The name database 10 is further characterized by an indication of the relationships between the different identifiers in separate classes. For example, FIG. 10 illustrates arrows 60 that link the different existing identifiers for a similar product. The direction of the arrow 60 in FIG. 10 shows a horizontal pattern used for hierarchical databases. However, arrow 60 may travel in any direction, in accordance with the possible relationships among the data in the name database 10.

An illustrative example is provided below:

EXAMPLE 1: A NAME DATABASE, LINKING TO INFORMATION FOUND AT
SEVERAL DIFFERENT SOURCES

<u>1A: Manufacturer's Database</u>	
MODEL	COLOR
r	RED
b	BLUE

<u>1B: Vendor 1 Database</u>	
SKU 1	COST
10	\$2

<u>1C: Vendor 2 Database</u>	
SKU 2	COST
100	\$3

<u>1D: Naming Database</u>		
MODEL	SKU 1	SKU 2
r	10	
b		100

In this example, the manufacturer produces two models, model r that is red and model b that is blue. However, the manufacturer does not provide information on the prices of the models. Vendor 1 sells a model with a SKU of 10 for \$2 and Vendor 2 sells a model with a SKU of 100 for \$3. However, neither Vendor 1 nor Vendor 2 indicates which model corresponds to the SKU employed by the vendor. Only through accessing the naming database can a consumer recognize that

Vendor 1 sells model r and Vendor 2 sells model b. In this way, the naming database serves as a modern Rosetta stone to associate the proprietary nomenclature from one source of product information with another source.

5 In the embodiment demonstrated in Example 1, the name database includes no information on the products, but instead only provides the identifiers and their interrelationships. It should be appreciated however, that the naming database could also include product information,
10 as seen in the following example.

<u>EXAMPLE 2: PRODUCT NAMES AND PRODUCT INFORMATION ARE ON THE SAME DATABASE</u>				
MODEL	SKU 1	SKU 2	COLOR	COST
r	10		RED	\$2
b		100	YELLOW	\$3
g	20	200	GREEN	\$3

15 In this example, the name database has combined the databases of Example 1, and information on a new model g is provided. As a result, the illustrated hierarchical database provides all known information on models r, b, and g. New model g, as indicated in the database, has a green color, costs \$3 and is available as SKU 20 at vendor 1 and as SKU 200 at vendor 2. In this example, new types of information are added to the database as additional columns
20 and additional products are added as new rows. In this example, as well as in Example 1, the relationships between the product identifiers are defined by the rows 30 and columns 20. In particular, different identifiers for the same product appear in the same row 30, and identifiers for
25

different products from the same source appear in the same column 40.

In addition, FIG. 10 illustrates product information columns 40 in the name database 10. As described above, the product information database 10 may include virtually any type of data related to the product. For example, the product information columns 40 may contain links to third party reviews of the particular product or to an Internet discussion regarding the product. Conversely, the product information may provide information on similar, competing products or indicate possible vendors for purchase to the product. The product information may further include related advertisements or pictures of the product.

As seen in the Cost column of Example 2, data entries may be redundant in a hierarchical base. To address this concern, the present invention preferably uses a relational database, as illustrated in the following example.

<u>EXAMPLE 3: RELATIONAL NAME DATABASE OF THE INFORMATION IN</u>					
<u>EXAMPLE 2</u>					
	MODEL	SKU 1	SKU 2	COLOR	COST
1	r	10	100	RED	\$2
2	b	20	200	YELLOW	\$3
3	g			GREEN	

With this relational database, a vector in the form of [model, SKU 1, SKU 2, color, cost] shows the relative relationship between the data in each column, rather than merely looking horizontally. In this example, the relationship vectors are [1,1,0,1,1], [2,0,1,2,2], and [3,2,2,3,2]. In other words, [1,1,0,1,1], corresponds to the first model (r), which has the first listed value of SKU

1(10), no value of SKU 2, the first listed color (red) and the first listed cost (\$2).

It should be appreciated that other database formations are possible and are well known in the field. The database structures illustrated in FIG 1 and the above examples may be easily modified to form different structures that perform the same function. For example, the name database 10 may be restructured so that new rows contain new data types and new columns contain additional members of known data types.

Similarly, the name database 10 may be multi-dimensioned. For instance, the name database 10 may have three dimensions: one to store the different products; a second to store the different names for the same product; and a third to store the various data about the product.

In one embodiment, name database 10 assigns a universal SKU 50 to every product. The universal SKU 50 may be, for example, an alphanumeric code. In this way, the name database 10 has a system for labeling the various products, which does not have to be altered as changes are made to the identifiers for the product. In another embodiment, the name database 10 is formed using SQL to permit easy additions and changes to the name database 10.

In order to make use of the normalized and associated information that is stored in the database, it must be capable of being queried by clients and presented or displayed in a readily understandable format. Queries against a standard relational database unfortunately do not perform satisfactorily to accommodate a large number of simultaneous clients (as is typically experienced by a website server), or to present a sophisticated user interface or display, even for a small number of users. Consequently, according to another aspect of the present

invention a product information server is provided which enables the information to be traversed and compared with query terms quickly.

According to this aspect of the invention, the object information is compiled into a compact, flat file format. The compact file format takes each character string for each piece of information and "tokenizes" it by assigning to it a unique integer. Although it is possible that the token may be arbitrarily chosen, according to the preferred embodiment of the invention the value of the integer assigned to the character string is equal to the offset of the location of the string in the data block. In this way, each token points to the beginning of its corresponding character string in the block. Consequently, the server is able to go immediately to the location of the start of the character string in the block based on the value of the token, so as to retrieve the string for display.

The character strings and unique integer values are placed in a look-up table 4100 as shown in Fig. 4B. Each character string is stored in a field 4102 which is associated with a unique integer value field 4101. In the example, the integer 2 identifies the character string "Pentium®", while the character string "CPU" is identified by integer 6598. Each of the tokens representing each product in the database is then written into a file 4001 having a format as shown in Fig. 4A.

Conventionally, information to be presented to a user in a table format is arranged in a file in product sequence order, with each product name being followed by all of the attribute data associated with the product. When organized into a table format, each row represents a specific product, each column represents a specific attribute of the product,

and each intersection of row and column contains a token for a character string corresponding to the attribute value.

Such a file is sometimes referred to as being in "row major" format. When carrying out a parameter search on such a

5 file, a great deal of irrelevant information is retrieved from the database (usually on a hard disk) and placed into memory. This has the double negative effect of using up the memory resources of the system and making the search take longer because of the need to scan through irrelevant

10 information. For example, if a search is desired for laptop computers having a minimum amount of memory, according to the conventional database file format all attribute information is retrieved for all laptop products, in addition to the attribute search term specified. Thus, the

15 search requires a substantial amount of time because all the irrelevant attribute information pertaining to each product in the database must be traversed in the course of identifying the pertinent attribute information specified by the user.

20 According to the invention, instead of arranging information in "row major" format, the product information server extracts the information from the native database and organizes it in "column major" format, wherein all attribute values of like attributes are arranged in sequence adjacent to each other. For example, all monitor display sizes are
25 arranged next to each other, then all display resolutions arranged next to each other, then all hard disk sizes are arranged next to each other, then all processor clock speeds are arranged next to each other, etc. In this way, an
30 attribute-based search may be performed much faster, by allowing the search to jump immediately to the start of the location of the relevant attribute specified by the user,

and to retrieve all the relevant attribute information and only the relevant attribute information into memory to perform the search.

As shown in Fig. 4A, N PROD 4003 is an integer identifying the number of products in the file, N ATTR 4005 is an integer identifying the number N of attributes in the file. Each of the N attributes is represented by an attribute value integer "ATTR I mval" 4007. The integer 4007 identifies the attribute. Each of the values in turn are identified by the "val I prod I" integers 4009. Additionally, an attribute may be multivalued, such that the integers 4007 would correspond to an offset for an "mval list I" 4013, which is an n-tuple, each of the n integers in the n-tuple pointing to a separate value of the attribute in the look-up table.

In a query, the file 4001 is traversed and all corresponding integers are retrieved. The associated character strings are then obtained from the look-up table 4100 and are appropriately formatted for display at the client.

As shown in FIG. 5, the present invention provides a system 400 to implement the method of the invention to achieve the desired information display. In particular, system 400 comprises a server 410 that contains a storage device 420 for storing the desired vendor and product information. The server also contains a database engine 425 that adds collected information data to the storage device 420 and creates an output using the information stored in the storage device 420.

The system 400 further includes a user's processing device 450, such as a personal computer, and a connection 440 to allow the transfer of information between the server

410 and the processing device 450. The processing device 450 includes a web browser 460 which provides an output to a display device 480, such as a display monitor, and which accepts an input from an input device 470, such as a keyboard or mouse.

In addition to the storage device 420, the server 410 also optionally contains scraper programs 430 for the collection of data, as previously described.

The connection 440 is preferably a distributed network, such as the Internet, to allow a plurality of users to have simultaneous connection to the server.

FIG. 6 illustrates a screen shot of a website containing information on a product specified by a user as being of interest and vendors that sell that product. The website displays a name 10 for the product, a list price 30, a composite user rating 40 based upon user ratings 45 in various categories 46, a ranking 50 of the product in a class 55 of similar products, features 60 of the product, vendors 70 who sell the product, a price 80 for the product at each of the vendors' sites, user reviews 90, and access to industry reviews 100.

The name 10 is generally the manufacturer and model name but may be any identifier used for the product. The name 10 may be carried over from a third-party site or arbitrarily created at the website.

Similarly, the list price 30 is a number either given by the product's manufacturer or distributor or arbitrarily assigned by the website. The list price 30 alerts a user to the relative value of the product to allow better evaluation of the prices 80 offered by the vendors 70. For instance, a computer selling for \$500 is generally a good value if its list price is \$1000, but not if the list price is \$100.

While the list price is generally higher than the actual price offered 80 by the vendors, this is not necessarily true, especially with rare, collectable items that may sell for much more than the list price.

5 The consumer product rating is formed, as described above, by surveying a plurality of users and combining these ratings.

10 As illustrated in FIG. 6, some of the vendors 70 may be identified prominently, so as to encourage the user to patronize these vendors. As further illustrated in FIG. 6, the website may optionally display any of the following: an image 20 of the product; a rate-it-now display 110 to allow the user to add a user review 90 and rating 40 of the product; a helpfulness evaluation 120 of the information; 15 complementary products 130 that may be purchased along with the desired product; or a discussion link 140 to Usenet and/or other discussion areas regarding the product and/or related products.

20 Because of limitations on the size of the display, the website may not all display of the product and vendor information at the same time. The information is then nested, and the consumer may access this information by performing an action such as clicking a pointing device (mouse) over one of the displayed objects. For example, to 25 find more information about one of the vendors 70, the user selects the vendor to be redirected to a sub-page, as shown in FIG. 7A. The sub-page then provides more specific information for the vendor 70, such as the vendor's address 71; telephone number 72; shipping practices 73; payment 30 policy 74; return policy 75; a rating of the vendor 76; reviews of the vendor 77; and an indication 78 of the product name 10, product prices 80, and availability 150.

The website may allow the user to select a product by reviewing a list of product categories 180, as illustrated in FIG. 7B. One the user selects a category of products, the user may then select a particular product from a product list 190 from that class, as shown in Fig. 7C. Alternatively, the product list 190 may be formed by displaying the highest rated products 170.

As illustrated in FIG. 8, the website may further contain a decision guide 300 which asks the user general questions 310 such as the user's age, occupation, and hobbies. The decision guide then uses this information to select a product for the user. This feature is helpful for a user who may not have sufficient technical knowledge to select a product based upon the features of that product. In this way, the product list 190 is formed to meet the specific needs of the user.

For a user who understands the product features, the website may assist the user in identifying products containing user-desired features. A narrow-your-choices option 160 of FIG. 6 redirects the user to a display, such as illustrated in FIG. 9. The narrow-your-choices option 160 asks the user to specify or select one or more feature options 161 for the product of interest. After the user has selected the desired feature options 161, the user sends a "display products" instruction 162 to the website to display the products meeting the chosen feature options 161. In this way, the product list 190 can be formed with products having the desired features.

The invention thus having been described, it will be apparent to those skilled in the art that the same may be varied in many ways without departing from the spirit and scope of the invention. Any and all such modifications are

intended to be included within the scope of the following claims.